

**IN THE CLAIMS:**

1. (Currently Amended) In a file server having a storage operating system, a method for  
2 managing storage of data in a plurality of storage devices, each storage device having a  
3 plurality of blocks for storing data, comprising:  
4 generating block layout information in a file system layer of the storage operating  
5 system by determining which blocks within the plurality of blocks are allocated for  
6 storing data and which are unallocated;  
7 transferring the block layout information from the file system layer to a RAID  
8 layer of the storage operating system; **and**  
9 responsive to the block layout information, controlling ~~the~~ execution of I/O  
10 operations at the RAID layer by identifying a plurality of contiguous blocks on a single  
11 storage device within the plurality of blocks for use by each I/O operation so as to  
12 substantially maximize chain lengths of read operations for calculation of parity;  
13 determining whether a parity subtraction method or a recalculation method  
14 requires a fewest number of read operations to calculate parity for the I/O operations;  
15 selecting the parity subtraction method or the recalculation method for parity  
16 calculation based on the which method that requires the fewest number of read  
17 operations to compute parity for the I/O operations; and  
18 responsive to the block layout information and the parity calculation method  
19 selected, identifying the contiguous blocks within the plurality of blocks for use by the  
20 I/O operations.

1. 2. (Currently Amended) A method for managing storage of data in a plurality of storage  
2 devices, each comprising a plurality of storage blocks, comprising:  
3 generating block layout information;  
4 determining whether a first methodology or a second methodology requires a  
5 fewest number of read operations to calculate parity; and  
6 in response to the block layout information and the determination, controlling  
7 execution of I/O operations by identifying a plurality of contiguous storage blocks on a  
8 single storage device for use by each I/O operation so as to substantially minimize the

9     number of read operations needed for calculation of error correction parameters across a  
10    | stripe disposed among the plurality of storage devices.

1     3. (Original) The method of claim 2 wherein the calculation of error correction  
2    parameters comprises the calculation of parity.

1     4. (Currently Amended) The method of claim 3-2 wherein the calculation of parity  
2    comprises selecting a parity calculation operation from a group consisting of a  
3    | subtraction method as the first methodology and a parity re-calculation method as the  
4    | second methodology.

1     5. (Currently Amended) The method of claim 2 wherein the identification of contiguous  
2    storage blocks for use in the I/O operation substantially maximizes a chain length by  
3    | substantially maximizing the-a number of blocks having a contiguous physical layout on  
4    | the single storage device.

1     6. (Currently Amended) The method of claim 2, further comprising:  
2     | identifying the contiguous storage blocks for use in the I/O operation so as to  
3    | substantially maximize the chain length by substantially maximizing the-a number of  
4    | blocks having sequential volume block numbers (VBNs) associated with the plurality of  
5    | storage blocks.

1     7. (Currently Amended) The method of claim 2, further comprising:  
2     | identifying the contiguous storage blocks for use in the I/O operation so as to  
3    | substantially maximize the chain length by substantially maximizing locality of the  
4    | contiguous blocks of the single storage device.

1     8. (Currently Amended) The method of claim 2 wherein the-controlling execution step  
2    comprises:

3 examining storage blocks to which data is to be written prior to write operations;  
4 and  
5 selecting one of a plurality of parity calculation methodologies including either,  
6 a-the first methodology comprising minimizing a number of blocks read, and-or  
7 a-the second methodology comprising maximizing chain lengths of blocks read  
8 for calculating parity calculation.

1 9. (Currently Amended) The method of claim 8, wherein the-controlling execution step  
2 further comprises:

3 implementing selection of the parity calculation methodology responsive to the  
4 block layout information; and

5 wherein, if the selection constitutes substantially-minimizing the number of  
6 blocks read,

7 determining on a stripe-by-stripe basis whether to calculate parity based on a  
8 subtraction method or a recalculation method,

9 performing any appropriate read operations to support the method selected, and  
10 calculating parity responsive to the number of blocks read and the data to be  
11 written; and

12 wherein, if the selection constitutes substantially-maximizing chain lengths of  
13 blocks read,

14 deciding which storage blocks to read to substantially-maximize chain length  
15 while substantially-minimizing the number of storage blocks read to support either a-the  
16 subtraction method or a-the recalculation method,

17 performing read operations on the number of blocks to-be-read, and  
18 calculating parity responsive to the number of blocks read and the data to be  
19 written.

1 10. (Currently Amended) The method of claim 2, wherein the identification of identifying  
2 the contiguous storage blocks is based at least in part on an available resource.

1 11. (Original) The method of claim 2 further comprising transmitting the block layout  
2 information from a file system layer to a RAID layer.

1 12. (Currently Amended) The method of claim 2 wherein ~~the generating step~~ further  
2 comprises:

3 making a first determination as to whether a storage block is unallocated;  
4 making a second determination as to a current implementation of the plurality of  
5 storage devices; and  
6 generating the block layout information based at least in part on the first and the  
7 second determinations.

1 13. (Currently Amended) The method of claim 2, wherein the I/O operation is one of a  
2 plurality of I/O operations and at least one of the plurality of I/O operations is a read  
3 operation.

1 14. (Currently Amended) The method of claim 2, wherein the chain length is a ~~chain~~  
2 length of a read operation for calculation of parity.

1 15. (Currently Amended) The method of claim 2, wherein the chain length is a ~~chain~~  
2 length for a write operation for the data.

1 16. (Currently Amended) A method for managing storage of data in a storage system,  
2 comprising:

3 maintaining a plurality of storage devices each having a plurality of storage  
4 blocks; and  
5 writing data to predetermined storage blocks of the plurality of storage blocks  
6 across a plurality of stripes and to predetermined contiguous storage blocks within each  
7 storage device so as to substantially maximize chain lengths of the predetermined  
8 contiguous storage blocks within each storage device and minimizing a number of read  
9 operations for the calculation of error correction parameters across each stripe of the

10 | plurality of stripes by determining whether a parity subtraction method or a recalculation  
11 | method requires a fewest number of read operations to calculate parity, and selecting thea  
12 | parity subtraction method or thea recalculation method for parity calculation based on the  
13 | which method that-requires the fewest number of read operations-to compute parity.

1 | 17. - 38. (Cancelled)

1 | 39. (Currently Amended) A storage system, comprising:  
2 | a plurality of storage devices each having a plurality of storage blocks; and  
3 | a storage manager in communication with the plurality of storage devices, the  
4 | storage manager configured to write writing data to predetermined storage blocks across  
5 | a plurality of stripes and to predetermined storage blocks within each storage device so as  
6 | to substantially maximize chain length of the plurality of storage blocks by selecting as  
7 | many as-contiguous storage blocks within a single storage device while substantially  
8 | minimizing a number of read operations required for calculation of error correction  
9 | parameters across each stripe of the plurality of stripes by determining whether a parity  
10 | subtraction method or a recalculation method requires a fewest number of read operations  
11 | to calculate parity and selecting a the parity subtraction method or a the recalculation  
12 | method for parity calculation based on the which method that-requires the fewest number  
13 | of read operations-to compute parity.

1 | 40. (Currently Amended) A system for managing the storage of data, the system  
2 | comprising:  
3 | a plurality of storage devices each having a plurality of storage blocks;  
4 | a storage device manager in communication with the plurality of storage blocks;  
5 | a block layout information generator in communication with the storage device  
6 | manager and the plurality of storage blocks; and  
7 | an error correction parameter calculator in communication with the plurality of  
8 | storage blocks and the storage device manager,

9       wherein the storage device manager, in response to the block layout information  
10      from the block layout information generator, controls the execution of an I/O operation  
11      by identifying a plurality of contiguous storage blocks on a single storage device for use  
12      by the I/O operation so as to substantially maximize chain length within the single  
13      storage device while substantially minimizing the a number of read operations required  
14      for calculation by the error correction parameter calculator of error correction parameters  
15      across a stripe by determining whether a parity subtraction method or a recalculation  
16      method requires a fewest number of read operations to calculate parity for the I/O  
17      operations and selecting a-the parity subtraction method or athe recalculation method for  
18      parity calculation based on the which method that requires the fewest number of read  
19      operations to compute parity.

1      41 - 44. (Cancelled)

1      45. (Currently Amended) A method for managing storage of data by a server,  
2      comprising:  
3      receiving a request to write the data to a plurality of storage devices;  
4      generating block layout information to determine which blocks within a plurality  
5      of blocks located in the plurality of storage devices are allocated for storing data and  
6      which are unallocated;  
7      identifying blocks within the plurality of blocks for use by a set of I/O operations  
8      to store the data;  
9      determining the a number of read operations needed to compute calculate parity  
10     for the data by computing calculating parity using a subtraction method-of computing  
11     parity;  
12     determining the a number of read operations needed to compute calculate parity  
13     for the data by computing calculating parity using a recalculation method-of computing  
14     parity;  
15     choosing either the subtraction method of computing calculating parity or the  
16     recalculation method of computing calculating parity by determining which of these two

17 | methods requires ~~the-a~~ fewer number of read operations, and choosing the which method  
18 | that requires -requiring the fewer number of read operations; and  
19 | writing the data to the identified blocks, and computing calculating parity for the  
20 | data using the chosen which method of computing parity.

1 | 46. (Currently Amended) The method of claim 45, further comprising:  
2 | choosing to either ~~firstly~~ maximize chain lengths of read operations for  
3 | ~~calculation of calculating~~ parity or ~~secondly~~ choosing to place the data with a high degree  
4 | of locality in the plurality of storage devices, by choosing the which method ~~which~~ that  
5 | requires the ~~fewest~~ fewer number of read operations ~~in~~ computing parity for the data.

1 | 47. (Currently Amended) A method for managing storage of data by a server,  
2 | comprising:  
3 | receiving a request to write data to a plurality of storage devices;  
4 | generating block layout information to determine which blocks within a plurality  
5 | of blocks located in the plurality of storage devices are allocated ~~for storing~~ data and  
6 | which are unallocated;  
7 | identifying blocks within the plurality of blocks for use by a set of I/O operations  
8 | to store the data; and

9 | selecting determining, in response to the block layout information, whether to  
10 | substantially minimize ~~the-a~~ number of read blocks or whether to substantially maximize  
11 | chain lengths of read blocks based on which method requires a fewer number of read  
12 | operations, and

13 | implementing ~~the-a~~ selection, during the writing of the data to the plurality of  
14 | storage devices, responsive to the block layout information, and responsive to whether  
15 | substantially minimize the number of read blocks or substantially maximize chain  
16 | lengths of read blocks requires the fewer number of read operations.

1 48. (Currently Amended) The method of claim 47, further comprising:  
2       in response to selecting to substantially-minimize the number of read blocks,  
3       determining whether to calculate parity based on the-a subtraction method or the-a  
4       recalculation method by determining based on which method requires the fewer number  
5       of read operations; and selecting the method which requires the fewer number of read  
6       operations; and  
7       performing the write operation and calculating the parity using the-which parity  
8       calculation method requiring requires the fewer number of read operations.

1 49. (Currently Amended) The method of claim-48-47, further comprising:  
2       in response to selecting to substantially-maximize chain lengths of read blocks,  
3       deciding which storage blocks to read to substantially-maximize chain lengths while  
4       minimizing the number of storage blocks read to support either the subtraction method or  
5       the recalculation method-of parity calculation; and  
6       performing the write operation and calculating the parity using the-which parity  
7       calculation method requiring requires the fewer number of read operations.

1 50. (Currently Amended) A method for managing storage of data by a server,  
2       comprising:  
3       receiving a request to write data to a plurality of storage devices;  
4       generating block layout information to determine which blocks within a plurality  
5       of blocks located in the plurality of storage devices are allocated for storing data and  
6       which are unallocated;  
7       identifying blocks within the plurality of blocks for use by a set of I/O operations  
8       to store the data;  
9       testing to either maximize chain lengths of read operations for calculation of  
10      parity, or to place the data with a high degree of locality in the plurality of storage  
11      devices, the testing having the steps comprising,  
12      determining, for both maximizing chain length and placing the data with a-the  
13      high degree of locality, the-a number of read operations needed to compute calculate

14 parity for the data, by ~~computing calculating~~ parity using both ~~the-a~~ subtraction method  
15 of ~~computing calculating~~ parity and ~~the-a~~ recalculation method of ~~computing calculating~~  
16 parity;  
17 firstly ~~first~~ choosing to either maximize chain lengths of read operations for  
18 calculation of parity or to place the data with ~~a~~~~the~~ high degree of locality in the plurality  
19 of storage devices, and after ~~this~~~~the~~ first choice, secondly choosing either the subtraction  
20 method of ~~computing calculating~~ parity or the recalculation method of ~~computing~~  
21 ~~calculating~~ parity by determining which of these methods requires ~~the-a~~ fewest number of  
22 read operations,  
23 choosing ~~the~~~~which~~ method ~~requiring~~~~requires~~ the fewest number of read  
24 operations of ~~computing calculating~~ parity of the data; and  
25 writing the data to ~~the~~ identified blocks, and ~~computing calculating~~ parity for the  
26 data using the ~~chosen~~~~which~~ method of ~~computing~~ parity.

1 51. (Currently Amended) A computer-readable media, comprising:  
2 said computer readable media containing instructions for execution on a processor  
3 for a method of managing storage of data in a plurality of storage devices, each storage  
4 device having a plurality of blocks for storing data, the method ~~having comprising~~,  
5 generating block layout information; and  
6 in response to the block layout information, controlling ~~the~~ execution of an I/O  
7 operation by identifying a plurality of contiguous storage blocks on a single storage  
8 device for use by the I/O operation so as to ~~substantially~~~~minimize~~ ~~the-a~~ number of read  
9 operations needed for calculation of error correction parameters across a stripe by  
10 ~~determining whether a parity subtraction method or a recalculation method requires a~~  
11 ~~fewest number of read operations to calculate parity for the I/O operations and selecting~~  
12 ~~the~~ parity subtraction method or ~~the~~ recalculation method for parity calculation based  
13 on the ~~which~~ method ~~that~~ requires the fewest number of read operations ~~to compute~~  
14 parity.